Summit Guard

AI PILOT GOVERNANCE KIT

Define the pilot decision

AI governance your team can use. Before your first pilot, record the decision, the boundaries and the evidence still needed.

Use one worksheet set for one proposed task. An assistant helps a person produce or review material. An agent may also call tools or take actions. Scope the authority as well as the task.

This is a planning aid for Australian business leaders. It is not a completed assessment or permission to launch. Examples are hypothetical; adjust them to your own context.

A proposed rule is a plan. An approved document records a decision. Evidence from a trial may show how a control behaved under the tested conditions. Keep these distinct.

Task and proposed tool

Business benefit to test and how you will measure it

Decision owner and people who need to be involved

Pilot limits: team, duration, sample size and systems

Summit Guard

AI PILOT GOVERNANCE KIT

Record tasks and information

Describe what the pilot will do before deciding whether the proposed information and permissions fit. Use information categories in this worksheet; do not add passwords, client records or personal details.

Separate public, internal, confidential and personal information. Record where inputs come from, what the tool may retain and where outputs go. Verify the actual account type, settings and vendor terms rather than inferring them from a product name.

Mark an unanswered question as unknown. Name the person who will resolve it before access expands.

Users, task and expected outputs

Input sources and permitted information categories

Information excluded; retention and sharing questions

Unknowns, evidence needed and responsible person

Plan assistant use

For a staff assistant, agree the allowed task, permitted inputs, output checks and escalation route. Keep a person accountable for decisions and communications.

Illustrative example: a small team drafts general website wording using public product information in an approved business account. A named person checks accuracy, unsupported claims and suitability before publication.

The example does not approve a particular product or its default settings. A task involving confidential proposals or personal information needs a fresh decision.

A useful review rule names the person, what they check, when they check it and what happens if they cannot resolve a concern.

Allowed tasks, tools and input conditions

Reviewer, review criteria and point before use or publication

Escalation contact and pause conditions

Trial examples to test and evidence to retain

Summit Guard

AI PILOT GOVERNANCE KIT

Plan agent authority

For an action-taking agent, record each permitted action, the affected system, the proposed permission boundary and the human approval point. Ask whether a less autonomous workflow could achieve the same benefit.

Illustrative example: an agent drafts follow-up messages from a test dataset. Sending remains a separate human action. A named owner can pause the trial; changes to sending permissions require a new decision.

Document an owner, distinct identity, access expiry, action records and a stop procedure. The implementation team must test that the actual permissions, approvals and stop mechanism enforce the intended limits.

A written stop procedure is not proof that an agent can be stopped. Avoid treating a documented design as technical validation.

Allowed actions, prohibited actions and affected systems

Proposed identity, permission limits and expiry

Human approvals; owner and authority to stop

Technical tests, records and recovery steps still needed

Summit Guard

AI PILOT GOVERNANCE KIT

Check the evidence

Use these prompts to prepare the evidence, not to calculate a maturity rating. Ask which statement each item supports and what it cannot establish.

Reported practice: an interview says staff will review outputs. Documented intention: the pilot plan assigns a reviewer. Evidence of operation: a controlled test records what was checked, by whom and with what result. Before a pilot, operation may be untested.

Useful items may include an approved use-case description, redacted vendor settings, a responsibility record, a review checklist and results from a controlled test. Request only the minimum needed.

Keep the original evidence in your own approved records. Record references here, rather than copying sensitive material. Missing evidence is a question to resolve, not an automatic pass or failure.

Claim or condition you need to support

Available evidence reference and date

What remains reported, planned or untested

Next evidence check, owner and due date

Summit Guard

AI PILOT GOVERNANCE KIT

Record the decision and next step

Allow: the accountable decision-maker accepts the bounded use under recorded conditions. Review: more evidence or a specialist decision is needed. Stop: pause the proposed use until the stated issue is resolved. These are decisions for your organisation.

Record conditions and a review date. Revisit the decision if the task, information, tool, permissions or affected people change, or if a trial raises a concern.

Need help defining the decision? Request a free 20-minute fit call at summitguard.com.au/contact. The AI Governance Risk Snapshot reviews up to three tools or use cases in one team and up to five agreed, redacted evidence items. Scope and fee are confirmed before work begins.

The call establishes suitability; it is not an assessment. The kit and Snapshot are not legal advice, certification, formal assurance or launch sign-off. Technical validation, implementation and ongoing monitoring are outside the Snapshot.

Reference points: National AI Centre essential AI practices (ai.gov.au/staying-safe-and-responsible/essential-ai-practices) and NIST AI RMF (nist.gov/itl/ai-risk-management-framework). Original Summit Guard worksheets; no maturity score or third-party endorsement. Version 1.0, 30 October 2026.

Decision: allow / review / stop; conditions and reason

Accountable decision-maker, date and evidence references

Open actions, owners and review date